arXiv:1606.02542v1 [cs.SD] 8 Jun 2016

Symbolic Music Data Version 1.0

Christian Walder
CSIRO Data61
7 London Circuit, Canberra, 2604, Australia.
christian.walder@data61.csiro.au

June 9, 2016

Abstract

In this document, we introduce a new dataset designed for training machine learning models of symbolic music data. Five datasets are provided, one of which is from a newly collected corpus of 20K midi files. We describe our preprocessing and cleaning pipeline, which includes the exclusion of a number of files based on scores from a previously developed probabilistic machine learning model. We also define training, testing and validation splits for the new dataset, based on a clustering scheme which we also describe. Some simple histograms are included.

1 Introduction

In this appendix we provide an overview of the symbolic music datasets we offer in pre-processed form[1]. Note that the source of four out of five of these datasets is the same set of midi files used in [BLBV12], which also provides pre-processed data. That work provided "piano roll" representations, which essentially consist of a regular temporal grid (of period one eighth note) of on/off indicators for each midi note number. While the piano roll is an excellent simplified music format for early investigations into symbolic music modelling, it does have several limitations, as discussed in previous work [Wal16]. To name one such limitation, the piano roll format does not explicitly represent note endings, and therefore cannot differentiate between, say, two successive eighth notes, and a single quarter note.

To address these limitations, we have extracted additional information from the same set of midi files. Our goal is to represent the performance (or sounding) of notes by when they begin and end, rather than whether they are sounding or not at each time on a regular temporal grid. The representation we adopt consists of sets of five-tuples of integers representing the:

- piece number (corresponding to a midi file),

[1]The data is available for download here: http://bit.ly/1PqNTJ2

Dataset	Long Name	Source	Total Pieces	Midi Resolution
PMD	`piano-midi.de`	[PE07, BLBV12]	124	480
JSB	J.S Bach Chorales	[AW05, BLBV12]	382	100
MUS	MuseData	[mus, BLBV12]	783	240
NOT	Nottingham	[not, BLBV12]	1037	480
CMA	Classical Midi Archives	[cla] (new)	19700	variable

Table 1: A summary of the datasets used in this study.

- track (or part) number, defined by the midi channel in which the note event occurs,
- midi note number, ranging 0-127 according to the midi standard, and 16-110 inclusive for the data we consider here,
- note start time, in "ticks", (2400 ticks = 1 beat = one quarter note),
- note end time, also in ticks.

This document provides some background on the data, with a special emphasis on our new relatively large dataset, which we derived from an archive kindly provided to us by Pierre Schwob of http://www.classicalarchives.com. We are permitted to release this data in the form we provide, but not to provide the original midi files. Please refer to the data archive itself[1] for a detailed description of the format.

A summary of the five datasets is provided in Table 1.

2 Preprocessing

We applied the following processing steps and filters to the raw midi data.

- Combination of piano "sustain pedal" signals with key press information to obtain equivalent individual note on/off events.
- Removal of duplicate/overlapping notes which occur on the same midi channel (while not technically allowed, this still occurs in real midi data due to the permissive nature of the midi file format). Unfortunately, this step is ill posed, and different midi software packages handle this differently. Our approach involves processing notes sequentially in order of start time, and ignoring those note events that overlap a previously added note event.
- Removal of midi channels with less than two note events (these occurred in the MUS and CMA datasets, and were always information tracks containing authorship information and acknowledgements, *etc.*).

- Removal of percussion tracks. These occurred in some of the Haydn symphonies and Bach Cantatas contained in the MUS dataset, as well as in the CMA dataset. It is important to filter these as the percussion instruments are not necessarily pitched, and hence the midi numbers in these tracks are not comparable with those of pitched instruments, which we aim to model.

- Re-sampling of the timing information to a resolution of 2400 ticks per quarter note, as this is the lowest common multiple of the original midi file resolutions (see Table 1) for the four datasets considered in [BLBV12]. We accept some quantization error for some of the CMA files, although 2400 is already an unusually fine grained midi quantization (*cf.* the resolutions of the other datasets, in Table 1).

For our new CMA dataset, we also removed 306 of the 20,006 midi files due to their suspect nature. We did this by assigning a heuristic score to each file and ranking. The score was computed by first training our model [Wal16] on the union of the four (transposed) datasets, JSB, PMD, NOT and MUS. We then computed the negative log-probability of each midi note number in the raw CMA data under the aforementioned model. Finally, we defined our heuristic score as, for each file, the mean of these negative log probabilities plus the standard error. The scores we obtained in this way are depicted in Figure 1. A listening test on the best and worst files verified the effectiveness of this measure. In any case, some degree of noise is to be expected in a data set of this size, and should be handled by subsequent modelling efforts.

3 Splits

The four datasets used in [BLBV12] retain the original training, testing, and validation splits used in that work. For CMA, we took a careful approach to data splitting. The main issue was data duplicates, since the midi archive we were provided contained multiple versions of several pieces, each encoded slightly differently by a different transcriber. To reduce the statistical dependence between the train/test/validation splits of the CMA set, we used the following methodology:

1. We computed a simple signature vector for each file, which consisted of the concatenation of two vectors. The first was the normalised histogram of midi note numbers in the file. For the second vector, we quantized the event durations into a set of sensible bins, and computed a normalised histogram of the resulting quantised durations.

2. Given the signature vectors associated with each file, we performed hierarchical clustering using the function `scipy.cluster.hierarchy.dendrogram` from the python scipy library[2]. We then ordered the files by traversing the resulting hierarchy in a depth first fashion.

[2]https://www.scipy.org

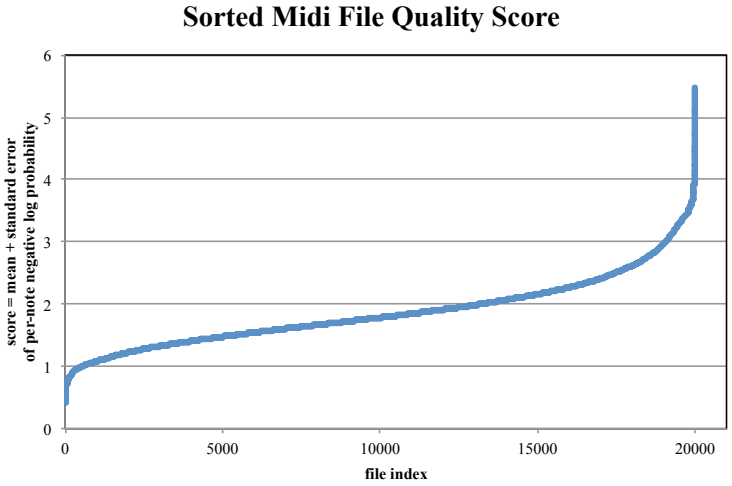

Figure 1: Our filtering score for the original 20,006 midi files provided by the website `http://www.classicalarchives.com`. We kept the top 19,700, discarding files with a score greater than 3.9.

3. Given the above ordering, we took contiguous chunks of 15,760, 1,970 and 1,970 files for the train, test, validation sets, respectively. This leads to a similar ratio of split sizes as in [BLBV12].

4 Basic Exploratory Plots

We provide some basic exploratory plots in figures 2–5.

The **Note Distribution** and **Number of Notes Per Piece** plots are self explanatory.

Note that the **Number of Parts Per Piece** (lower left sub figure) is fixed at one for the entire JSB dataset. This is due to an unfortunate lack of midi track information in those files, many of which are in fact four part harmonies. The pieces in the NOT dataset feature either one part (in the case of pure melodies) or two (in the case of melodies with associated chord accompaniments). The PMD dataset features up to six parts (for a three-part Bach fugue in which left and right hands are tracked separately). MUS features up to 27 parts (for Bach's *St. Matthew's Passion*). The CMA data features two pieces with 46 parts — Ravel's *Valses Nobles et Sentimentales*, and *Venus*, by Gustav Holst.

The least obvious sub-figures are those on the lower-right labeled **Peak Polyphonicity Per Piece**. Polyphonicity simply refers to the number of simultaneously sounding notes, and this number can be rather high. For the PMD data, this is mainly attributable to musical "runs" which are performed with the piano sustain pedal depressed, for example in some of the Liszt pieces.

For the MUS data, this is mainly due to the inclusion of large orchestral works which feature many instruments. The CMA data, of course, contains both of the aforementioned sources of high levels of polyphonicity.

Acknowledgements

Special thanks to Pierre Schwob of http://www.classicalarchives.com, who permitted us to release the data in the form we describe.

References

[AW05] M. Allan and C. K. I. Williams, *Harmonising Chorales by Probabilistic Inference*, Advances in Neural Information Processing Systems 17 (2005).

[BLBV12] N. Boulanger-Lewandowski, Y. Bengio and P. Vincent, Modeling Temporal Dependencies in High-Dimensional Sequences: Application to Polyphonic Music Generation and Transcription, in *Proceedings of the Twenty-ninth International Conference on Machine Learning (ICML'12)*, ACM, 2012.

[cla] `www.classicalarchives.com`.

[mus] `www.musedata.org`.

[not] `ifdo.ca/~seymour/nottingham/nottingham.html`.

[PE07] G. E. Poliner and D. P. W. Ellis, *A Discriminative Model for Polyphonic Piano Transcription.*, EURASIP J. Adv. Sig. Proc. **2007** (2007).

[Wal16] C. Walder, *Modelling Symbolic Music: Beyond the Piano Roll*, arXiv (2016), 1606.01368.

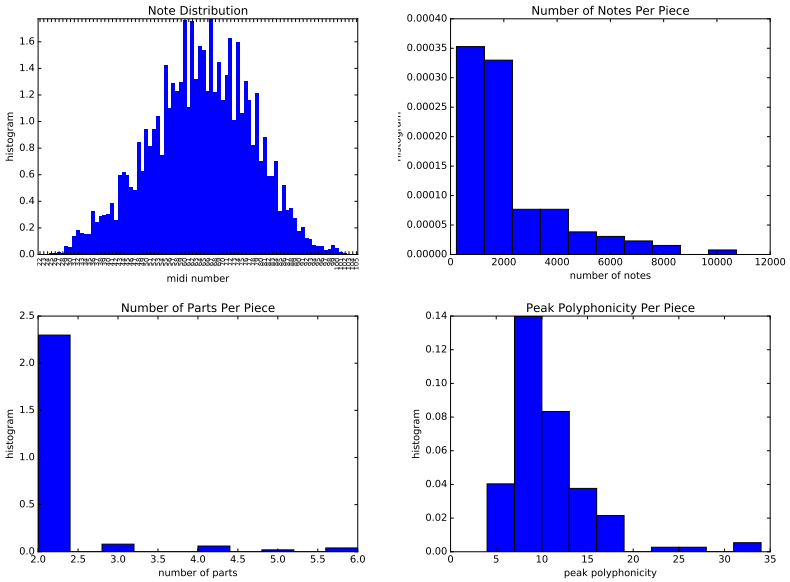

Figure 2: Summary of the PMD dataset.

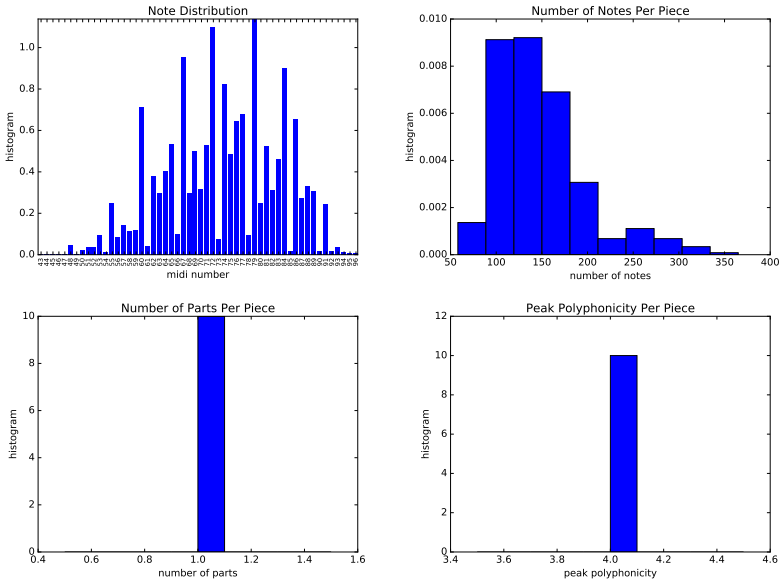

Figure 3: Summary of the JSB dataset.

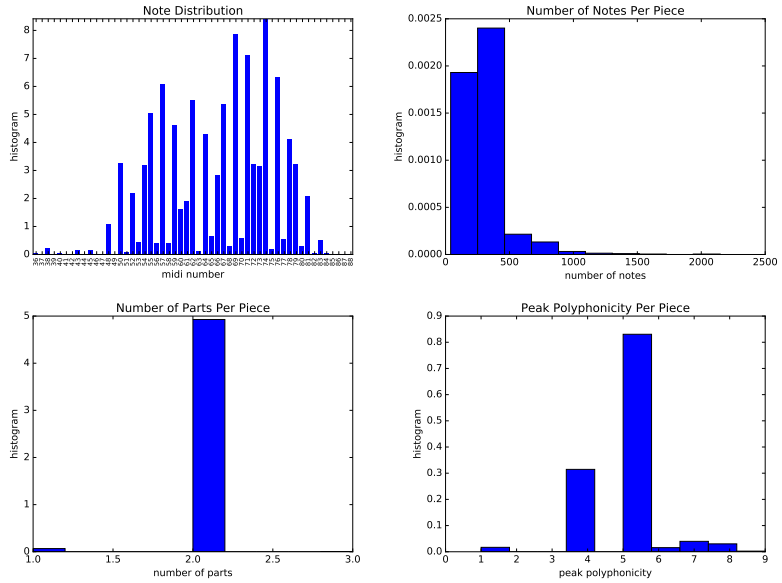

Figure 4: Summary of the NOT dataset.

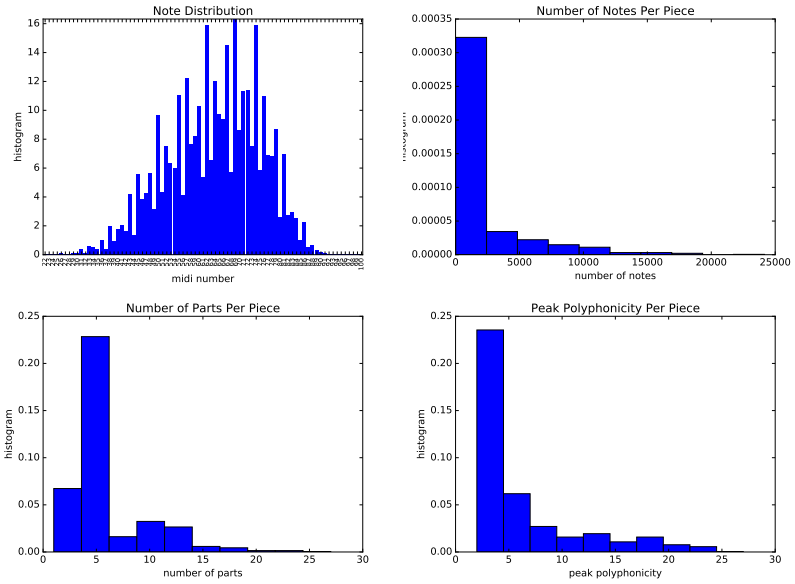

Figure 5: Summary of the MUS dataset.

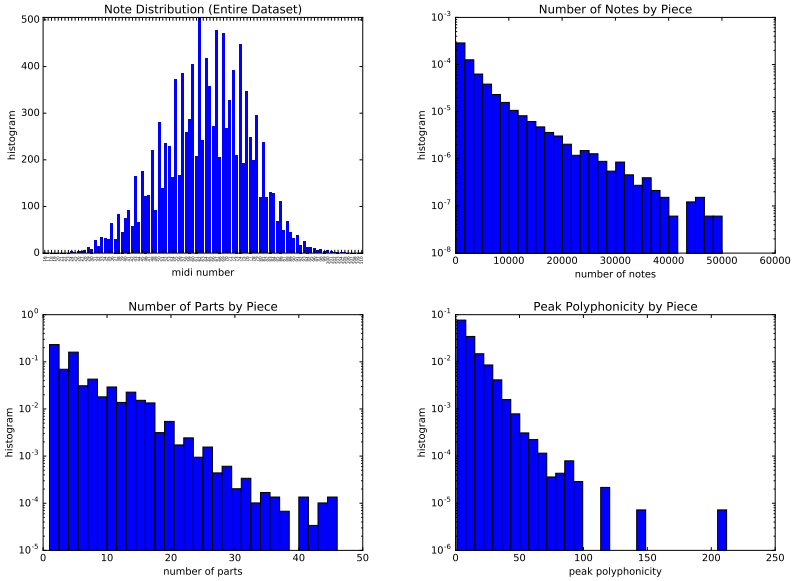

Figure 6: Summary of the CMA dataset. Note the log scale on three of the plots.